

Linear Convergence of Accelerated Stochastic Gradient Descent for Nonconvex Nonsmooth Optimization

Feihu Huang

HUANGFEIHU@NUAA.EDU.CN

College of Computer Science and Technology

Nanjing University of Aeronautics and Astronautics, Nanjing, 210016, China

Songcan Chen

S.CHEN@NUAA.EDU.CN

College of Computer Science and Technology

Nanjing University of Aeronautics and Astronautics, Nanjing, 210016, China

Editor:

Abstract

In this paper, we study the stochastic gradient descent (SGD) method for the nonconvex nonsmooth optimization, and propose an accelerated SGD method by combining the variance reduction technique with Nesterov's extrapolation technique. Moreover, based on the local error bound condition, we establish the linear convergence of our method to obtain a stationary point of the nonconvex optimization. In particular, we prove that not only the sequence generated linearly converges to a stationary point of the problem, but also the corresponding sequence of objective values is linearly convergent. Finally, some numerical experiments demonstrate the effectiveness of our method. To the best of our knowledge, it is first proved that the accelerated SGD method converges linearly to the local minimum of the nonconvex optimization.

Keywords: Linear convergence, Stochastic gradient descent, Variance reduction, Error bound, Nonconvex nonsmooth optimization

1. Introduction

Stochastic gradient descent (SGD) is a widely useful optimization method in machine learning, due to its simplicity and scalability (Bottou, 2004; Bousquet and Bottou, 2008). For example, the SGD only computes gradient of one sample instead of all samples in each iteration. However, due to the variance in the stochastic process, its scalability of the SGD is limited by its slower convergence rate. Recently, some accelerated versions of the SGD have successfully been proposed to reduce the variance, and obtain faster convergence rate than the standard SGD. For example, the stochastic average gradient (SAG) (Roux et al., 2012) obtains a linear convergence rate in solving the strongly convex problem by incorporating the old gradients estimated in the previous iterations. The stochastic dual coordinate ascent (SDCA) (Shalev-Shwartz and Zhang, 2013) performs the stochastic coordinate ascent on the dual problems, and also obtains a linear convergence rate in solving strongly convex problems. However, these algorithms have faster convergence rate than the standard SGD at the cost of requiring much space to store old gradients or dual variables. Thus, Johnson and Zhang (2013); Xiao and Zhang (2014) have proposed the stochastic variance reduced gradient (SVRG) algorithms via using a multi-stage scheme to progressively reduce the vari-

ance of the stochastic gradient, which obtain a fast convergence rate with no extra space to store the intermediate gradients or dual variables. At the same time, a novel method called SAGA has been proposed in Defazio et al. (2014), which extends the SAG method and enjoys better theoretical convergence rate than both SAG and SVRG. Moreover, Nitanda (2014) has proved an accelerated SGD method incorporating the variance reduction (VR) technique (Johnson and Zhang, 2013) and the Nesterov’s acceleration method (Nesterov, 2004), and obtains a faster convergence rate than that of the SVRG.

So far, the above stochastic gradient methods rely mainly on strongly convex or convex objective functions. However, there exists many useful nonconvex models in machine learning such as some robust empirical risk minimization models (Aravkin and Davis, 2016) and deep learning (LeCun et al., 2015). More recently, some works begin focusing on the stochastic gradient methods for the nonconvex optimizations. For example, Ghadimi and Lan (2016); Ghadimi et al. (2016) have established the iteration complexity of $O(1/\epsilon^2)$ for the standard SGD to obtain an ϵ -stationary solution. Allen-Zhu and Hazan (2016); Reddi et al. (2016a) have proved that the variance reduced SGD methods reach the iteration complexity of $O(1/\epsilon)$ for solving the nonconvex problems. Moreover, Reddi et al. (2016b); Aravkin and Davis (2016) have studied the variance reduced SGD methods for solving the nonconvex nonsmooth problems, and have proved that these methods have the iteration complexity of $O(1/\epsilon)$ to reach an ϵ -stationary point.




Figure 1: The relationships between some conditions: *invex functions*, *convex*, *strongly convex*, *PL condition* and *error bound condition*. Note that a function is invex if and only if every stationary point is a global minimum.

In particular, Reddi et al. (2016b) has established the linear convergence of the proposed SGD methods for the objective functions that satisfy the Polyak-Lojasiewicz (PL) inequality (Polyak, 1963; Karimi et al., 2016). In fact, when the objective functions satisfy the PL condition, every stationary point of the problem is a global minimum Polyak (1963); Karimi et al. (2016). At the same time, Aravkin and Davis (2016) also has established the linear convergence of the proposed SGD methods for the nonconvex nonsmooth problems, based on a modified globalization of the error bound, which is stricter than the standard error bound (Luo and Tseng, 1993; Tseng, 2010). In fact, this global error bound condition is equivalent to combining the standard error bound condition with the invex function, whose every stationary point is a global minimum Ben-Israel and Mond (1986). For the sake of clarity, Figure 1 shows the relationships between some conditions such as the invex functions, PL condition and the error bound condition. In summary, Reddi et al. (2016b); Aravkin and Davis (2016) only prove the linear convergence of their methods for the nonconvex

optimization, which has a global minimum. Due to that there rarely exists the global minimum in the nonconvex problems, their convergence analysis only satisfies few nonconvex problems. At present, thus, there still exists an open question:

“How to design a novel stochastic gradient method with a linear convergence rate for solving the general nonconvex nonsmooth problems without a global minimum?”

In the paper, thus, we try to design a new stochastic gradient method for solving the nonconvex nonsmooth problems, and establish the linear convergence of our method, based on the local error bound, which is much looser than the PL condition used in Reddi et al. (2016b) and the global error bound condition used in Aravkin and Davis (2016). Specifically, we propose an accelerated SGD method by incorporating the VR technique and the Nesterov’s extrapolation technique, to minimize the nonconvex nonsmooth problems as follows:

$$\min_{x \in R^p} \frac{1}{n} \sum_{i=1}^n f_i(x) + g(x), \quad (1)$$

where $f(x) = \frac{1}{n} \sum_{i=1}^n f_i(x)$ is the sum of a set of nonconvex functions that have Lipschitz continuous gradients; $g(x)$ is a lower semi-continuous convex function that is non-smooth. The problem (1) is inspired by the regularized empirical risk minimization in machine learning Vapnik (2013). For example, $f(x)$ denotes the loss function such as the sigmoid loss, and $g(x)$ denotes the regularization term such as sparse regularization $\|x\|_1$. Throughout the paper, let $\Phi(x) = f(x) + g(x)$. In summary, our main contributions are three-fold as follows:

- We propose an accelerated SGD method for the non-convex non-smooth optimization by incorporating the VR technique and Nesterov’s extrapolation technique.
- We establish the linear convergence of our method to obtain a stationary point of the problem, based on the local error bound. In particular, we prove that not only the generated sequence R -linearly converges to a stationary point, but also the corresponding sequence of objective values is R -linearly convergent.
- Some numerical experiments support the effectiveness of our method.

1.1 Notations and Preliminaries

$\|\cdot\|$ denotes the Euclidean norm; $\|\cdot\|_1$ and $\|\cdot\|_\infty$ denote the ℓ_1 norm and ℓ_∞ norm, respectively. For a nonempty closed set $\mathcal{C} \subseteq R^p$, $\text{dist}(x, \mathcal{C}) = \inf_{y \in \mathcal{C}} \|x - y\|$ denotes the distance from x to \mathcal{C} . $\text{dom}(f)$ denotes the domain of function f .

Next, we recall two notions of (local) linear convergence rate. For a sequence $\{x_t\}$, if there exist $c \in (0, 1)$ and $t_0 > 0$ such that

$$\|x_{t+1} - x_*\| \leq c \|x_t - x_*\|, \quad \forall t \geq t_0,$$

we say that it Q -linearly converges to x_* , and the sequence $\{x_t\}$ R -linearly converges to x_* , if

$$\limsup_{t \rightarrow \infty} \|x_t - x_*\|^{\frac{1}{t}} < 1.$$

In the following, we use the relation between the above notions of linear convergence rate.

Lemma 1 *Wen et al. (2017)* Suppose that $\{a_k\}$ and $\{b_k\}$ are two sequences in R with $0 \leq b_k \leq a_k$ for all k , and $\{a_k\}$ is Q -linearly convergent to zero. Then $\{b_k\}$ is R -linearly convergent to zero.

2. Accelerated SGD for Nonconvex Nonsmooth Optimization

In the section, we propose an accelerated SGD method for solving the nonconvex nonsmooth problems, by combining the VR and the Nesterov's extrapolation techniques.

Firstly, we observe that any smooth function $f(x)$ can be written as $f(x) = f_L(x) - f_l(x)$, where both $f_L(x)$ and $f_l(x)$ are convex and smooth functions. For example, the smooth function $f(x)$ can be decomposed as

$$f(x) = \underbrace{f(x) + \frac{c}{2}\|x\|^2}_{f_L(x)} - \underbrace{\frac{c}{2}\|x\|^2}_{f_l(x)}, \quad (2)$$

for any $c \geq L$, where L is the Lipschitz constant of $\nabla f(x)$. Clearly, $f_l(x)$ is a convex function, which has the Lipschitz continuous gradient with the constant c . The following Lemma proves that $f_L(x)$ is a convex and smooth function.

Lemma 2 *The function $f_L(x) = f(x) + \frac{c}{2}\|x\|^2$ in (2) is convex and smooth function.*

A detailed proof of Lemma 2 is provided in Appendix A. Without loss of generality, thus, we assume $f(x) = f_L(x) - f_l(x)$ for some convex functions $f_L(x)$ and $f_l(x)$ with Lipschitz continuous gradients with $L \geq l$, where L and l denote the Lipschitz constants of $\nabla f_L(x)$ and $\nabla f_l(x)$, respectively. It is not hard to show that $\nabla f(x)$ is Lipschitz continuous with the constant L .

Algorithm 1 Accelerated SGD (ASVRG) for Non-convex Non-smooth Optimization

- 1: **Input:** epoch length m , S , step size $\mu > 0$, and the constant $\sigma > \frac{1}{2}$;
 - 2: **Initialize:** $x_m^0 = \tilde{y}^0 = x^0 \in R^p$, $\{\beta_t\} \subseteq [0, \sqrt{\frac{2}{2+\sigma}})$;
 - 3: **for** $s = 0, 1, \dots, S-1$ **do**
 - 4: $x_0^{s+1} = x_{-1}^{s+1} = x_m^s$;
 - 5: $\nabla f(\tilde{y}^s) = \frac{1}{n} \sum_{i=1}^n \nabla f_i(\tilde{y}^s)$;
 - 6: **for** $t = 0, 1, \dots, m-1$ **do**
 - 7: Uniformly randomly pick $I_t \subseteq \{1, 2, \dots, n\}$ (with replacement) such that $|I_t| = b$;
 - 8: $y_t^{s+1} = x_t^{s+1} + \beta_t(x_t^{s+1} - x_{t-1}^{s+1})$;
 - 9: $v_t^{s+1} = \frac{1}{b} \sum_{i_t \in I_t} (\nabla f_{i_t}(y_t^{s+1}) - \nabla f_{i_t}(\tilde{y}^s)) + \nabla f(\tilde{y}^s)$;
 - 10: $x_{t+1}^{s+1} = \text{Prox}_{\mu g}(y_t^{s+1} - \mu v_t^{s+1})$;
 - 11: **end for**
 - 12: $\tilde{y}^{s+1} = x_m^{s+1}$;
 - 13: **end for**
 - 14: **Output:** Iterate x chosen uniformly random from $\{(x_t^s)_{t=0}^m\}_{s=1}^S$.
-

Next, we propose an accelerated SGD algorithm for solving the problem (1), which is described in Algorithm 1. In Algorithm 1, we use the $\text{Prox}_g(v)$ to denote the proximal

operator of a proper closed convex function $g(x)$ at any $v \in R^p$, i.e.,

$$\text{Prox}_g(v) = \arg \min_{x \in R^p} \left\{ \frac{1}{2} \|x - v\|^2 + g(x) \right\}. \quad (3)$$

So the update x in the step 10 of Algorithm 1 is equivalently given by

$$x_{t+1}^{s+1} = \arg \min_{x \in R^p} \left\{ \langle v_t^{s+1}, x \rangle + \frac{1}{2\mu} \|x - y_t^{s+1}\|^2 + g(x) \right\}. \quad (4)$$

Since $\mathbb{E}[v_t^{s+1}] = \mathbb{E}[\frac{1}{b} \sum_{i_t \in I_t} (\nabla f_{i_t}(y_t^{s+1}) - \nabla f_{i_t}(\tilde{y}^s)) + \nabla f(\tilde{y}^s)] = \nabla f(y_t^{s+1})$, the stochastic gradient v_t^{s+1} is unbiased.

In Algorithm 1, we let $y_t^{s+1} = x_t^{s+1} + \beta_t(x_t^{s+1} - x_{t-1}^{s+1})$, which adopts the Nesterov's acceleration technique. Moreover, due to using the variance reduction technique as the SVRG in Johnson and Zhang (2013), not only the stochastic gradient v_t^{s+1} is unbiased, but also its variance is progressively reduced. In the following, we give an upper bound of the variance of the stochastic gradient v_t^{s+1} .

Lemma 3 *In Algorithm 1, we have $v_t^{s+1} = \frac{1}{b} \sum_{i_t \in I_t} (\nabla f_{i_t}(y_t^{s+1}) - \nabla f_{i_t}(\tilde{y}^s)) + \nabla f(\tilde{y}^s)$, $y_t^{s+1} = x_t^{s+1} + \beta_t(x_t^{s+1} - x_{t-1}^{s+1})$ and $\tilde{y}^s = x_m^s$, and set $\Delta_t^{s+1} = v_t^{s+1} - \nabla f(y_t^{s+1})$, then the following inequality holds*

$$\mathbb{E}\|\Delta_t^{s+1}\|^2 \leq \frac{L^2}{2b} \|x_t^{s+1} - x_m^s\|^2 + \frac{\bar{\beta}^2 L^2}{2b} \|x_t^{s+1} - x_{t-1}^{s+1}\|^2, \quad (5)$$

where $\bar{\beta} = \max_t \beta_t$, and $\mathbb{E}\|\Delta_t^{s+1}\|^2$ denotes the variance of the stochastic gradient v_t^{s+1} .

A detailed proof of Lemma 3 is provided in Appendix B. Lemma 3 shows that the variance of the stochastic gradient v_t^{s+1} has a upper bound $O(\|x_t^{s+1} - x_{t-1}^{s+1}\|^2 + \|x_t^{s+1} - x_m^s\|^2)$. As number of iterations increases, all x_t^{s+1} , x_{t-1}^{s+1} and x_m^s approach the same stationary point, thus the variance of stochastic gradient vanishes.

3. Convergence Analysis

In this section, we study the convergence behavior of our method. Based on the local error bound condition, we prove that our method has a local linear convergence rate to obtain a stationary point. First, we give some mild assumptions as follows:

Assumption 1 *For $\forall i \in \{1, 2, \dots, n\}$, the gradient of function f_i is Lipschitz continuous with the constant $L_i > 0$, such that for $x_1, x_2 \in R^d$*

$$\|\nabla f_i(x_1) - \nabla f_i(x_2)\| \leq L_i \|x_1 - x_2\| \leq L \|x_1 - x_2\|, \quad (6)$$

where $L = \max_i L_i$, and this is equivalent to

$$f_i(x_1) \leq f_i(x_2) + \nabla f_i(x_2)^T (x_1 - x_2) + \frac{L}{2} \|x_1 - x_2\|^2. \quad (7)$$

Assumption 2 *The function $\Phi(x)$ is level bounded, i.e., the set $\{x \in R^p : \Phi(x) \leq r\}$ is bounded for any $r \in R$, and it is bounded below, i.e., $\inf_x \Phi(x) > -\infty$.*

Assumption 3 Let \mathcal{X} denote the set of stationary points of $\Phi(x)$.

(i) For any $\xi \geq \inf_{x \in \mathbb{R}^p} \Phi(x)$, there exist $\epsilon > 0$ and $\gamma > 0$, such that

$$\text{dist}(x, \mathcal{X}) \leq \gamma \mathbb{E} \|\text{Prox}_{\mu g}(x - \mu \nabla f(x)) - x\|, \quad (8)$$

whenever $\mathbb{E} \|\text{Prox}_{\mu g}(x - \mu \nabla f(x)) - x\| \leq \epsilon$ and $\Phi(x) \leq \xi$.

(ii) There exists $\delta > 0$, such that $\|x - y\| \geq \delta$ whenever $x, y \in \mathcal{X}$, $\Phi(x) \neq \Phi(y)$.

Assumption 1 is widely used in the analysis of proximal gradient descent methods (Wen et al., 2017) and stochastic optimization methods (Allen-Zhu and Hazan, 2016; Huang et al., 2016). Assumption 2 gives the bound of objective function, which also is widely used in the analysis of proximal gradient decent methods (Wen et al., 2017). Assumption 3 gives the local error bound condition, which has been used in the convergence analysis of many algorithms (Tseng, 2010; Wen et al., 2017; Zhou and So, 2015; Drusvyatskiy and Lewis, 2016), such as the block coordinate gradient descent and proximal gradient descent method.

In Assumption 3, due to the solution obtained by using stochastic gradient descent method, so we use $\mathbb{E} \|\text{Prox}_{\mu g}(x - \mu \nabla f(x)) - x\|$ to replace $\|\text{Prox}_{\mu g}(x - \mu \nabla f(x)) - x\|$ in the local error bound. While, Aravkin and Davis (2016) uses a modified globalization of the error bound to the convergence analysis as follows:

$$\Phi(x) - \Phi(x_*) \leq \gamma \|\text{Prox}_{\mu g}(x - \mu \nabla f(x)) - x\|^2, \quad (9)$$

where $x_* = \arg \min_x \Phi(x)$ is an optimal solution of (1). Clearly, this global error bound (9) is much stricter than the local error bound, due to the fact that there are few global solution in the nonconvex optimizations.

In the following, we will give some useful lemmas for the convergence analysis of our algorithm. First, assume that the sequence $\{(x_t^s)_{t=0}^m\}_{s=1}^S$ is generated from Algorithm 1, then we define an useful sequence $\{(H_t^s)_{t=0}^m\}_{s=1}^S$ as follows:

$$H_t^s = \mathbb{E} [\Phi(x_t^s) + \alpha (\|x_t^s - x_m^{s-1}\|^2 + \|x_t^s - x_{t-1}^s\|^2)], \quad (10)$$

where α is a positive constant.

Lemma 4 Assume that the sequence $\{(x_t^s)_{t=0}^m\}_{s=1}^S$ is generated from Algorithm 1. Given the parameters $\{\beta_t\} \subseteq [0, \sqrt{\frac{2}{2+\sigma}})$ with $\sigma > 0$, if the constant α satisfies

$$\alpha = \frac{2\sigma}{2\sigma - 1} \left((1 + \bar{\beta}^2) \frac{L^2}{4\tau b} + \bar{\beta}^2 \left(\frac{1}{2\mu} + \frac{l}{2} \right) \right) > 0, \quad (11)$$

where $\bar{\beta} = \max_t \beta_t$, and the step size μ satisfies the following inequality

$$\mu \leq \min \left\{ \frac{1}{L}, \frac{2 - (\sigma + 2)\bar{\beta}^2}{2\tau + (\sigma + 2) \left(\frac{(\bar{\beta}^2 + 1)L^2}{2\tau b} + \bar{\beta}^2 l \right)} \right\}, \quad (12)$$

where the constant $0 < \tau < L$, then the following statements hold:

(i) For any $z \in \text{dom}(g)$, we have

$$\begin{aligned} \mathbb{E}[\Phi(x_{t+1}^{s+1})] &\leq \Phi(z) + \frac{L^2}{4\tau b} \|x_t^{s+1} - x_m^s\|^2 + \frac{\bar{\beta}^2 L^2}{4\tau b} \|x_t^{s+1} - x_{t-1}^{s+1}\|^2 \\ &\quad + \left(\frac{1}{2\mu} + \frac{l}{2}\right) \|z - y_t^{s+1}\|^2; \end{aligned} \quad (13)$$

(ii) For all $s > 0$ and any $t \in \{0, 1, \dots, m-1\}$

$$H_{t+1}^{s+1} - H_t^{s+1} \leq -\kappa(\|x_{t+1}^{s+1} - x_t^{s+1}\|^2 + \|x_t^{s+1} - x_{t-1}^{s+1}\|^2 + \|x_t^{s+1} - x_m^s\|^2), \quad (14)$$

where $0 \leq \kappa = \min \left\{ \frac{1}{2\mu} - \frac{\tau}{2} - \frac{\sigma+2}{2}\alpha, \bar{\beta}^2 \left(\frac{L^2}{4\tau b} + \frac{1}{2\mu} + \frac{l}{2} \right), \frac{L^2}{4\tau b} + \frac{\alpha}{2\sigma} \right\}$;

(iii) The sequence $\{(H_t^s)_{t=0}^m\}_{s=1}^S$ is non-increasing.

A detailed proof of Lemma 4 is provided in Appendix C. Lemma 4 shows that the sequence $\{(H_t^s)_{t=0}^m\}_{s=1}^S$ is non-increasing, which is useful to the following analysis.

Lemma 5 Under the conditions of Lemma 4 and the Assumption 2, the following statements hold:

- (i) The sequence $\{(H_t^s)_{t=0}^m\}_{s=1}^\infty$ is convergent;
- (ii) $\sum_{s=0}^\infty \sum_{t=0}^m (\|x_t^{s+1} - x_m^s\|^2 + \|x_t^{s+1} - x_{t-1}^{s+1}\|^2) \leq \infty$;
- (iii) Any accumulation point of $\{(x_t^s)_{t=0}^m\}_{s=1}^\infty$ is a stationary point of $\Phi(x)$.

A detailed proof of Lemma 5 is provided in Appendix D. Lemma 5 shows that the sequence $\{(H_t^s)_{t=0}^m\}_{s=1}^S$ is convergent under some mild conditions.

Lemma 6 Let $\{(x_t^s)_{t=0}^m\}_{s=1}^\infty$ be a sequence generated from Algorithm 1, and its set of accumulation points be denoted by Ω . Then there exists $\lim_{s \rightarrow \infty} \mathbb{E}[\Phi(x_t^s)] = \zeta$ for any $t \in \{0, 1, \dots, m\}$, and $\Phi \equiv \zeta$ on Ω .

A detailed proof of Lemma 6 is provided in Appendix E. The Lemma 6 shows that the function $\Phi(x)$ is the constant on the accumulation points generated from Algorithm 1.

Theorem 7 Let $\{(x_t^s)_{t=0}^m\}_{s=1}^\infty$ be a sequence generated from Algorithm 1. Under the above Assumptions and the conditions in Lemma 4, the following statements hold:

- (i) $\lim_{s \rightarrow \infty} \text{dist}(x_t^s, \mathcal{X}) = 0$ for any $t \in \{0, 1, 2, \dots, m\}$;
- (ii) The sequence $\{(H_t^s)_{t=0}^m\}_{s=1}^\infty$ is Q -linearly convergent.

A detailed proof of Theorem 7 is provided in Appendix F. Theorem 7 shows that the sequence $\{(H_t^s)_{t=0}^m\}_{s=1}^\infty$ is Q -linearly convergent. Finally, based the above results, we will prove that both the sequences $\{(x_t^s)_{t=0}^m\}_{s=1}^\infty$ and $\{(\Phi(x_t^s))_{t=0}^m\}_{s=1}^\infty$ are R -linearly convergent in the following.

Theorem 8 *The above assumptions hold, and Let $\{(x_t^s)_{t=0}^m\}_{s=1}^\infty$ be a sequence generated by Algorithm 1. Then the following statements hold:*

- (i) *The sequence $\{(x_t^s)_{t=0}^m\}_{s=1}^\infty$ R -linearly converges to a stationary point of the problem (1);*
- (ii) *The sequence $\{(\Phi(x_t^s))_{t=0}^m\}_{s=1}^\infty$ is R -linearly convergent.*

Proof From Theorem 7, the sequence $\{(H_t^s)_{t=0}^m\}_{s=1}^\infty$ is Q -linearly convergent, and let $\lim_{s \rightarrow \infty} H_t^s = \zeta$. By the conclusion (ii) of Lemma 5, we have, for any $t \in \{0, 1, \dots, m\}$

$$\|x_t^{s+1} - x_m^s\|^2 \leq \frac{1}{\kappa}(H_t^s - \zeta) - \frac{1}{\kappa}(H_{t+1}^s - \zeta) \leq \frac{1}{\kappa}(H_t^s - \zeta), \quad (15)$$

where the last inequality holds by the fact that the sequence $\{(H_t^s)_{t=0}^m\}_{s=1}^\infty$ is non-increasing and convergent to ζ .

Using (15) and the fact that the sequence $\{(H_t^s)_{t=0}^m\}_{s=1}^\infty$ is Q -linearly convergent, there exist $M > 0$ and $0 < c < 1$ such that

$$\|x_t^{s+1} - x_m^s\| \leq Mc^s$$

for all s and any $t \in \{0, 1, \dots, m\}$. Thus, for any $k_2 > k_1 \geq 1$, we have

$$\|x_t^{k_2} - x_m^{k_1}\| \leq \sum_{s=k_1}^{k_2-1} \|x_t^{s+1} - x_m^s\| \leq \frac{Mc^{k_1}}{1-c}. \quad (16)$$

Since (16) implies that the sequence $\{(x_t^s)_{t=1}^m\}_{s=1}^\infty$ is a Cauchy sequence, it is convergent. Let \hat{x} denote the limit of the sequence $\{(x_t^s)_{t=1}^m\}_{s=1}^\infty$, then we have, for any $t \in \{0, 1, \dots, m\}$

$$\lim_{k_2 \rightarrow \infty} \|x_t^{k_2} - x_m^{k_1}\| = \|\hat{x} - x_m^{k_1}\| \leq \frac{Mc^{k_1}}{1-c}, \quad (17)$$

Thus, using the Lemma 6 to (17), we can conclude that the sequence $\{(x_t^s)_{t=0}^m\}_{s=1}^\infty$ is R -linearly convergent to the stationary point of $\Phi(x)$.

Next, we prove (ii). By the definition of the sequence $\{H_t^s\}$, we have, for all $s \geq 1$, any $t \in \{1, 2, \dots, m\}$

$$\begin{aligned} |\Phi(x_t^s) - \zeta| &= |H_t^s - \zeta - \alpha(\|x_t^s - x_m^{s-1}\|^2 + \|x_t^s - x_{t-1}^s\|^2)| \\ &\leq H_t^s - \zeta + \alpha(\|x_t^s - x_m^{s-1}\|^2 + \|x_t^s - x_{t-1}^s\|^2) \\ &\stackrel{(i)}{\leq} H_t^s - \zeta + \frac{\alpha}{\kappa}(\|x_t^s - x_m^{s-1}\|^2 + \|x_t^s - x_{t-1}^s\|^2), \end{aligned} \quad (18)$$

where the inequality (i) holds by (14). Using Lemma 1 to (18), we conclude that the sequence $\{(\Phi(x_t^s))_{t=0}^m\}_{s=1}^\infty$.

■

Table 1: Summary of data sets and regularization parameters used in our experiments.

data sets	number of samples	dimensionality	λ_1	λ_2
<i>a9a</i>	32,561	123	10^{-5}	1.2×10^{-5}
<i>coverttype</i>	581,012	54	10^{-5}	1.2×10^{-5}
<i>rcv1</i>	677,399	47,236	10^{-4}	1.5×10^{-4}

4. Numerical Experiments

In this section, we conduct some numerical experiments to demonstrate the effectiveness of our algorithm. In the experiments, we mainly focus on the binary classification problem with nonconvex loss function, and compare our algorithm (ASVRG) with the nonconvex SVRG and SAGA in Reddi et al. (2016b).




Figure 2: Non-convex binary classification. Performances of ASVRG, SVRG and SAGA on the 'a9a', 'coverttype' and 'rcv1' datasets. Here, the y -axis is the solution suboptimality i.e., $\|x - x_*$, where x_* denotes the best solution obtained by running the proximal gradient descent with multiple restarts.




Figure 3: Non-convex binary classification. Performances of ASVRG, SVRG and SAGA on the 'a9a', 'coverttype' and 'rcv1' datasets. Here, the y -axis is the function suboptimality i.e., $f(x) - f(x_*)$, where x_* denotes the best solution obtained by running the proximal gradient descent with multiple restarts.

Specifically, given a set of straining samples $\{(a_i, b_i)\}_{i=1}^n$, where $a_i \in R^p$, $b_i \in \{-1, +1\}$, $\forall i \in \{1, 2, \dots, n\}$, then we solve the nonconvex problem as follows:

$$\min_x \frac{1}{n} \sum_{i=1}^n f_i(x) + \lambda_1 \|x\|_1 + \lambda_2 \|x\|_2^2, \quad (19)$$

where $f_i(x) = \frac{1}{1+\exp(b_i a_i^T x)}$ is a nonconvex sigmoid loss function (Allen-Zhu and Hazan, 2016; Huang et al., 2016). Here, λ_1 and λ_2 are the positive regularization parameters, respectively.

In the experiments, we use standard machine learning datasets in LIBSVM¹, which are summarized in Table 1. In addition, we use a constant minibatch size $b = 1$ in the algorithms. For the ASVRG and SVRG, we use the epoch length $m = n$. Finally, we report the objective function values and solutions for the datasets. Specifically, we report the function suboptimality, i.e., $f(x_t^{s+1}) - f(x_*)$ (for ASVRG and SVRG) and $f(x_t) - f(x_*)$ (for SAGA), and the solution suboptimality, i.e., $\|x_t^{s+1} - x_*\|$ (for ASVRG and SVRG) and $\|x_t - x_*\|$ (for SAGA). Here, x_* represents the best solution obtained by running proximal gradient descent with multiple restarts. In the algorithms, we compare both the solution and function suboptimality for the number of effective passes through the datasets, where each effective pass estimates n component gradients.

Figure 2 shows that the solution of ASVRG is closer to the best solution x_* compared to both the SVRG and SAGA. This implies that our algorithm has faster convergence than both the SVRG and SAGA. Moreover, Figure 3 shows that the objection function of ASVRG is much lower compared to both the SVRG and SAGA. This also implies that our algorithm has faster convergence than both the SVRG and SAGA.

5. Conclusion

In this paper, we have proposed an accelerated SGD method for solving the nonconvex nonsmooth problems via incorporating the VR and the Nesterov’s extrapolation techniques. Moreover, based on the local error bound condition, we prove that the generated sequence $\{x_t^s\}$ R -linearly converges to a stationary point, and the corresponding sequence of objective value $\{\Phi(x_t^s)\}$ also is R -linearly convergent. Finally, some numerical experiments demonstrate the effectiveness of our method.

Appendix A. The proof of Lemma 2

Proof Since the function $f(x)$ has the Lipschitz continuous gradient with a constant L , $-f(x)$ also has the Lipschitz continuous gradient with a constant L , thus we have

$$-f(x_2) \leq -f(x_1) + \langle -\nabla f(x_1), x_2 - x_1 \rangle + \frac{L}{2} \|x_2 - x_1\|^2. \quad (20)$$

By (20), we have

$$f(x_2) \geq f(x_1) + \langle \nabla f(x_1), x_2 - x_1 \rangle - \frac{L}{2} \|x_2 - x_1\|^2. \quad (21)$$

1. The datasets can be downloaded from <https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets>.

By $f_l(x) = \frac{c}{2}\|x\|^2$, we have

$$f_l(x_2) = f_l(x_1) + \langle \nabla f_l(x_1), x_2 - x_1 \rangle + \frac{c}{2}\|x_2 - x_1\|^2. \quad (22)$$

Summing (21) and (22), we have

$$f_L(x_2) \geq f_L(x_1) + \langle \nabla f_L(x_1), x_2 - x_1 \rangle + \frac{c-L}{2}\|x_2 - x_1\|^2.$$

From the above result and $c \geq L$, the function $f_L(x)$ is convex. Since both the functions $f(x)$ and $f_l(x)$ are smooth, the function $f_L(x) = f(x) + f_l(x)$ is smooth. \blacksquare

Appendix B. The proof of Lemma 3

Proof Since $v_t^{s+1} = \frac{1}{b} \sum_{i_t \in I_t} (\nabla f_{i_t}(y_t^{s+1}) - \nabla f_{i_t}(\tilde{y}^s)) + \nabla f(\tilde{y}^s)$, we have

$$\begin{aligned} & \mathbb{E}\|v_t^{s+1} - \nabla f(y_t^{s+1})\|^2 \\ &= \mathbb{E}\left\|\frac{1}{b} \sum_{i_t \in I_t} (\nabla f_{i_t}(y_t^{s+1}) - \nabla f_{i_t}(\tilde{y}^s) + \nabla f(\tilde{y}^s) - \nabla f(y_t^{s+1}))\right\|^2 \\ &= \frac{1}{b^2} \mathbb{E}\left\|\sum_{i_t \in I_t} (\nabla f_{i_t}(y_t^{s+1}) - \nabla f_{i_t}(\tilde{y}^s) + \nabla f(\tilde{y}^s) - \nabla f(y_t^{s+1}))\right\|^2 \\ &\leq \frac{1}{b} \mathbb{E}\|\nabla f_{i_t}(y_t^{s+1}) - \nabla f_{i_t}(\tilde{y}^s) + \nabla f(\tilde{y}^s) - \nabla f(y_t^{s+1})\|^2 \\ &\stackrel{(i)}{=} \frac{1}{b} (\mathbb{E}\|\nabla f_{i_t}(y_t^{s+1}) - \nabla f_{i_t}(\tilde{y}^s)\|^2 - \|\nabla f(y_t^{s+1}) - \nabla f(\tilde{y}^s)\|^2) \\ &\leq \frac{1}{b} \mathbb{E}\|\nabla f_{i_t}(y_t^{s+1}) - \nabla f_{i_t}(\tilde{y}^s)\|^2 \\ &= \frac{1}{nb} \sum_{i=1}^n \|\nabla f_{i_t}(y_t^{s+1}) - \nabla f_{i_t}(\tilde{y}^s)\|^2 \\ &\stackrel{(ii)}{\leq} \frac{L^2}{b} \|y_t^{s+1} - \tilde{y}^s\|^2, \end{aligned} \quad (23)$$

where the equality (i) holds by the equality $\mathbb{E}(\xi - \mathbb{E}\xi)^2 = \mathbb{E}\xi^2 - (\mathbb{E}\xi)^2$ for random variable ξ ; the inequality (ii) holds by Assumption 1. Next, using $y_t^{s+1} = x_t^{s+1} + \beta_t(x_t^{s+1} - x_{t-1}^{s+1})$ and $\tilde{y}^s = x_m^s$ in (23), we have

$$\begin{aligned} \mathbb{E}\|v_t^{s+1} - \nabla f(y_t^{s+1})\|^2 &\leq \frac{L^2}{b} \|y_t^{s+1} - \tilde{y}^s\|^2 \\ &= \frac{L^2}{b} \|x_t^{s+1} + \beta_t(x_t^{s+1} - x_{t-1}^{s+1}) - x_m^s\|^2 \\ &\stackrel{(i)}{\leq} \frac{L^2}{2b} \|x_t^{s+1} - x_m^s\|^2 + \frac{\bar{\beta}^2 L^2}{2b} \|x_t^{s+1} - x_{t-1}^{s+1}\|^2, \end{aligned}$$

where the inequality (i) holds by the inequality of $(a+b)^2 \leq 2a^2 + 2b^2$. \blacksquare

Appendix C. The proof of Lemma 4

Proof First, using the strong convexity of the objective in the minimization problem (4), for any $z \in \text{dom}(g)$, we have

$$g(x_{t+1}^{s+1}) \leq g(z) + \langle v_t^{s+1}, z - x_{t+1}^{s+1} \rangle + \frac{1}{2\mu} \|z - y_t^{s+1}\|^2 - \frac{1}{2\mu} \|x_{t+1}^{s+1} - y_t^{s+1}\|^2 - \frac{1}{2\mu} \|x_{t+1}^{s+1} - z\|^2. \quad (24)$$

Next, using $\nabla f(x)$ is Lipschitz continuous with the constant L , we have

$$f(x_{t+1}^{s+1}) \leq f(y_t^{s+1}) + \langle \nabla f(y_t^{s+1}), x_{t+1}^{s+1} - y_t^{s+1} \rangle + \frac{L}{2} \|x_{t+1}^{s+1} - y_t^{s+1}\|^2. \quad (25)$$

Summing (24) and (25), we have

$$\begin{aligned} f(x_{t+1}^{s+1}) + g(x_{t+1}^{s+1}) &\leq f(y_t^{s+1}) + g(z) + \langle v_t^{s+1} - \nabla f(y_t^{s+1}), z - x_{t+1}^{s+1} \rangle + \langle \nabla f(y_t^{s+1}), z - y_t^{s+1} \rangle \\ &\quad + \frac{1}{2\mu} \|z - y_t^{s+1}\|^2 - \left(\frac{1}{2\mu} - \frac{L}{2}\right) \|x_{t+1}^{s+1} - y_t^{s+1}\|^2 - \frac{1}{2\mu} \|x_{t+1}^{s+1} - z\|^2. \end{aligned} \quad (26)$$

Considering $f(x) = f_L(x) - f_l(x)$, we have

$$\begin{aligned} f(y_t^{s+1}) + \langle \nabla f(y_t^{s+1}), z - y_t^{s+1} \rangle &= f_L(y_t^{s+1}) - f_l(y_t^{s+1}) + \langle \nabla f_L(y_t^{s+1}), z - y_t^{s+1} \rangle \\ &\quad - \langle \nabla f_l(y_t^{s+1}), z - y_t^{s+1} \rangle. \end{aligned} \quad (27)$$

Since $f_L(x)$ and $f_l(x)$ are convex, and $\nabla f_L(x)$ and $\nabla f_l(x)$ are Lipschitz continuous with the constants L and l , respectively, we have

$$f_L(y_t^{s+1}) + \langle \nabla f_L(y_t^{s+1}), z - y_t^{s+1} \rangle \leq f_L(z), \quad (28)$$

$$f_l(z) - f_l(y_t^{s+1}) - \langle \nabla f_l(y_t^{s+1}), z - y_t^{s+1} \rangle \leq \frac{l}{2} \|z - y_t^{s+1}\|^2. \quad (29)$$

Combining (27), (28) with (29), we have

$$f(y_t^{s+1}) + \langle \nabla f(y_t^{s+1}), z - y_t^{s+1} \rangle \leq f(z) + \frac{l}{2} \|z - y_t^{s+1}\|^2. \quad (30)$$

Set $\Phi(x) = f(x) + g(x)$, and combining (26) and (30), we have

$$\begin{aligned} \Phi(x_{t+1}^{s+1}) &\leq \Phi(z) + \langle v_t^{s+1} - \nabla f(y_t^{s+1}), z - x_{t+1}^{s+1} \rangle + \left(\frac{1}{2\mu} + \frac{l}{2}\right) \|z - y_t^{s+1}\|^2 \\ &\quad - \left(\frac{1}{2\mu} - \frac{L}{2}\right) \|x_{t+1}^{s+1} - y_t^{s+1}\|^2 - \frac{1}{2\mu} \|z - x_{t+1}^{s+1}\|^2 \\ &\stackrel{(i)}{\leq} \Phi(z) + \frac{1}{2\tau} \|v_t^{s+1} - \nabla f(y_t^{s+1})\|^2 + \left(\frac{1}{2\mu} + \frac{l}{2}\right) \|z - y_t^{s+1}\|^2 \\ &\quad - \left(\frac{1}{2\mu} - \frac{L}{2}\right) \|x_{t+1}^{s+1} - y_t^{s+1}\|^2 - \left(\frac{1}{2\mu} - \frac{\tau}{2}\right) \|z - x_{t+1}^{s+1}\|^2, \end{aligned} \quad (31)$$

where the inequality (i) holds by the Cauchy inequality. Then taking expectation conditioned on information I_t to (31), we have

$$\begin{aligned}
\mathbb{E}[\Phi(x_{t+1}^{s+1})] &\leq \Phi(z) + \frac{1}{2\tau} \mathbb{E}[\|v_t^{s+1} - \nabla f(y_t^{s+1})\|^2] + (\frac{1}{2\mu} + \frac{l}{2}) \|z - y_t^{s+1}\|^2 \\
&\quad - (\frac{1}{2\mu} - \frac{L}{2}) \|x_{t+1}^{s+1} - y_t^{s+1}\|^2 - (\frac{1}{2\mu} - \frac{\tau}{2}) \|z - x_{t+1}^{s+1}\|^2 \\
&\stackrel{(i)}{\leq} \Phi(z) + \frac{L^2}{4\tau b} \|x_t^{s+1} - x_m^s\|^2 + \frac{\bar{\beta}^2 L^2}{4\tau b} \|x_t^{s+1} - x_{t-1}^{s+1}\|^2 + (\frac{1}{2\mu} + \frac{l}{2}) \|z - y_t^{s+1}\|^2 \\
&\quad - (\frac{1}{2\mu} - \frac{L}{2}) \|x_{t+1}^{s+1} - y_t^{s+1}\|^2 - (\frac{1}{2\mu} - \frac{\tau}{2}) \|z - x_{t+1}^{s+1}\|^2,
\end{aligned} \tag{32}$$

where the inequality (i) holds by Lemma 2, and $0 < \tau < L$. Since $\frac{1}{\mu} > L > \tau$, we can obtain the inequality (13). This proves (i).

Using (32) with $z = x_t^{s+1}$, we have

$$\begin{aligned}
\mathbb{E}[\Phi(x_{t+1}^{s+1})] &\leq \Phi(x_t^{s+1}) + \frac{L^2}{4\tau b} \|x_t^{s+1} - x_m^s\|^2 + \frac{\bar{\beta}^2 L^2}{4\tau b} \|x_t^{s+1} - x_{t-1}^{s+1}\|^2 + (\frac{1}{2\mu} + \frac{l}{2}) \|x_t^{s+1} - y_t^{s+1}\|^2 \\
&\quad - (\frac{1}{2\mu} - \frac{L}{2}) \|x_{t+1}^{s+1} - y_t^{s+1}\|^2 - (\frac{1}{2\mu} - \frac{\tau}{2}) \|x_{t+1}^{s+1} - x_t^{s+1}\|^2 \\
&\stackrel{(i)}{\leq} \Phi(x_t^{s+1}) + \frac{L^2}{4\tau b} \|x_t^{s+1} - x_m^s\|^2 + \bar{\beta}^2 (\frac{L^2}{4\tau b} + \frac{1}{2\mu} + \frac{l}{2}) \|x_t^{s+1} - x_{t-1}^{s+1}\|^2 \\
&\quad - (\frac{1}{2\mu} - \frac{L}{2}) \|x_{t+1}^{s+1} - y_t^{s+1}\|^2 - (\frac{1}{2\mu} - \frac{\tau}{2}) \|x_{t+1}^{s+1} - x_t^{s+1}\|^2,
\end{aligned} \tag{33}$$

where the inequality (i) holds by the definition of $y_t^{s+1} = x_t^{s+1} + \beta_t(x_t^{s+1} - x_{t-1}^{s+1})$. Next, considering $\|x_{t+1}^{s+1} - x_m^s\|^2$, we have

$$\begin{aligned}
\|x_{t+1}^{s+1} - x_m^s\|^2 &= \|x_{t+1}^{s+1} - x_t^{s+1} + x_t^{s+1} - x_m^s\|^2 \\
&\leq \frac{\sigma}{2} \|x_{t+1}^{s+1} - x_t^{s+1}\|^2 + \frac{1}{2\sigma} \|x_t^{s+1} - x_m^s\|^2,
\end{aligned} \tag{34}$$

where $\sigma > 0$. Then by combining (33) and (34), we have

$$\begin{aligned}
&\mathbb{E}[\Phi(x_{t+1}^{s+1}) + \alpha(\|x_{t+1}^{s+1} - x_m^s\|^2 + \|x_{t+1}^{s+1} - x_t^{s+1}\|^2)] \\
&\leq \Phi(x_t^{s+1}) + (\frac{\alpha}{2\sigma} + \frac{L^2}{4\tau b}) \|x_t^{s+1} - x_m^s\|^2 + \bar{\beta}^2 (\frac{L^2}{4\tau b} + \frac{1}{2\mu} + \frac{l}{2}) \|x_t^{s+1} - x_{t-1}^{s+1}\|^2 \\
&\quad - (\frac{1}{2\mu} - \frac{L}{2}) \|x_{t+1}^{s+1} - y_t^{s+1}\|^2 - (\frac{1}{2\mu} - \frac{\tau}{2} - \frac{\sigma+2}{2}\alpha) \|x_{t+1}^{s+1} - x_t^{s+1}\|^2 \\
&\stackrel{(i)}{\leq} \Phi(x_t^{s+1}) + \alpha(\|x_t^{s+1} - x_m^s\|^2 + \|x_t^{s+1} - x_{t-1}^{s+1}\|^2) \\
&\quad - (\frac{1}{2\mu} - \frac{L}{2}) \|x_{t+1}^{s+1} - y_t^{s+1}\|^2 - (\frac{1}{2\mu} - \frac{\tau}{2} - \frac{\sigma+2}{2}\alpha) \|x_{t+1}^{s+1} - x_t^{s+1}\|^2 \\
&\quad - \bar{\beta}^2 (\frac{L^2}{4\tau b} + \frac{1}{2\mu} + \frac{l}{2}) \|x_t^{s+1} - x_m^s\|^2 - (\frac{L^2}{4\tau b} + \frac{\alpha}{2\sigma}) \|x_t^{s+1} - x_{t-1}^{s+1}\|^2,
\end{aligned} \tag{35}$$

where the inequality (i) holds by (11) and (12).

Using (12), we have $\frac{1}{2\mu} - \frac{L}{2} \geq 0$ and $\frac{1}{2\mu} - \frac{\tau}{2} - \frac{\sigma+2}{2}\alpha \geq 0$. Thus, we can conclusion that $H_{t+1}^{s+1} \leq H_t^{s+1}$ for $t = 0, 1, \dots, m-1$. In addition, by $x_{-1}^{s+1} = x_0^{s+1} = x_m^s$, for any $s \in \{1, 2, \dots, S\}$

$$\begin{aligned} H_0^{s+1} &= \mathbb{E}[\Phi(x_0^{s+1}) + a(\|x_0^{s+1} - x_m^s\|^2 + \|x_0^{s+1} - x_{-1}^{s+1}\|^2)] \\ &= \Phi(x_m^s) \\ &\leq \mathbb{E}[\Phi(x_m^s) + a(\|x_m^s - x_m^{s-1}\|^2 + \|x_m^s - x_{m-1}^s\|^2)] \\ &= H_m^s. \end{aligned}$$

The sequence $\{(H_t^s)_{t=1}^m\}_{s=1}^S$ monotonically decreases over $t \in \{1, 2, \dots, m\}$ in each epoch $s \in \{1, 2, \dots, S\}$, and $H_m^s \geq H_0^{s+1}$ for any $s \in \{1, 2, \dots, S-1\}$, so it is non-increasing. This proves (iii).

Finally, using (35), we have

$$\begin{aligned} H_{t+1}^{s+1} &\leq H_t^{s+1} - \left(\frac{1}{2\mu} - \frac{L}{2}\right)\|x_{t+1}^{s+1} - y_t^{s+1}\|^2 - \left(\frac{1}{2\mu} - \frac{\tau}{2} - \frac{\sigma+2}{2}\alpha\right)\|x_{t+1}^{s+1} - x_t^{s+1}\|^2 \\ &\quad - \bar{\beta}^2\left(\frac{L^2}{4\tau b} + \frac{1}{2\mu} + \frac{l}{2}\right)\|x_t^{s+1} - x_m^s\|^2 - \left(\frac{L^2}{4\tau b} + \frac{\alpha}{2\sigma}\right)\|x_t^{s+1} - x_{t-1}^{s+1}\|^2 \\ &\stackrel{(i)}{\leq} H_t^{s+1} - \left(\frac{1}{2\mu} - \frac{\tau}{2} - \frac{\sigma+2}{2}\alpha\right)\|x_{t+1}^{s+1} - x_t^{s+1}\|^2 - \bar{\beta}^2\left(\frac{L^2}{4\tau b} + \frac{1}{2\mu} + \frac{l}{2}\right)\|x_t^{s+1} - x_m^s\|^2 \\ &\quad - \left(\frac{L^2}{4\tau b} + \frac{\alpha}{2\sigma}\right)\|x_t^{s+1} - x_{t-1}^{s+1}\|^2 \\ &\leq -\kappa(\|x_{t+1}^{s+1} - x_t^{s+1}\|^2 + \|x_t^{s+1} - x_{t-1}^{s+1}\|^2 + \|x_t^{s+1} - x_m^s\|), \end{aligned}$$

where the inequality (i) holds by the inequality (12), and $\kappa = \min\left\{\frac{1}{2\mu} - \frac{\tau}{2} - \frac{\sigma+2}{2}\alpha, \bar{\beta}^2\left(\frac{L^2}{4\tau b} + \frac{1}{2\mu} + \frac{l}{2}\right), \frac{L^2}{4\tau b} + \frac{\alpha}{2\sigma}\right\}$. Thus, this proves (ii). \blacksquare

Appendix D. The proof of Lemma 5

Proof From Lemma 4, the sequence $\{(H_t^s)_{t=0}^m\}_{s=1}^\infty$ is non-increasing. This together with the definition of H_t^s implies that

$$\Phi(x_t^s) \leq H_t^s \leq H_0^1 \leq \infty. \quad (36)$$

Since the function $\Phi(x)$ is level bound by Assumption 2, the sequence $\{(x_t^s)_{t=1}^m\}_{s=1}^\infty$ is bound. Moreover, using $\inf_x \Phi(x) > -\infty$ from Assumption 2, we have

$$H_t^s = \mathbb{E}[\Phi(x_t^s) + \alpha(\|x_t^s - x_m^{s-1}\|^2 + \|x_t^s - x_{t-1}^s\|^2)] > -\infty. \quad (37)$$

Thus, we can conclude that the sequence $\{(H_t^s)_{t=0}^m\}_{s=1}^\infty$ is convergent. This proves (i).

From the conclusion (ii) of Lemma 4, we have

$$\kappa(\|x_t^{s+1} - x_m^s\|^2 + \|x_t^{s+1} - x_{t-1}^{s+1}\|^2) \leq H_t^{s+1} - H_{t+1}^{s+1}. \quad (38)$$

Summing both sides of (38) on t from 0 to m and s from 0 to $S - 1$, we have

$$\sum_{s=0}^{S-1} \sum_{t=0}^m (\|x_t^{s+1} - x_m^s\|^2 + \|x_t^{s+1} - x_{t-1}^{s+1}\|^2) \leq \frac{1}{\kappa} (H_0^1 - H_m^S). \quad (39)$$

Since the sequence $\{(H_t^s)_{t=0}^m\}_{s=1}^\infty$ is convergent by the conclusion (i), and letting $S \rightarrow \infty$ in (39), we have

$$\sum_{s=0}^\infty \sum_{t=0}^m (\|x_t^{s+1} - x_m^s\|^2 + \|x_t^{s+1} - x_{t-1}^{s+1}\|^2) \leq \lim_{S \rightarrow \infty} \frac{1}{\kappa} (H_0^1 - H_m^S) < \infty. \quad (40)$$

This proves (ii).

Finally, we prove (iii). Let \bar{x} be an accumulation point, then there exists a subsequence $\{x_t^{s_i}\}$ such that $\lim_{i \rightarrow \infty} x_t^{s_i} = \bar{x}$ for any $t \in \{0, 1, \dots, m\}$. Using the first-order optimality condition of (4) and $v_t^{s_i+1} = \frac{1}{b} \sum_{i_t \in I_t} (\nabla f_{i_t}(y_t^{s_i+1}) - \nabla f_{i_t}(x_m^{s_i})) + \nabla f(x_m^{s_i})$, we have

$$-\frac{1}{\mu}(x_{t+1}^{s_i+1} - y_t^{s_i+1}) - \frac{1}{b} \sum_{i_t \in I_t} (\nabla f_{i_t}(y_t^{s_i+1}) - \nabla f_{i_t}(x_m^{s_i})) \in \nabla f(x_m^{s_i}) + \partial g(x_{t+1}^{s_i+1}), \quad (41)$$

where $y_t^{s_i+1} = x_t^{s_i+1} + \beta_t(x_t^{s_i+1} - x_{t-1}^{s_i+1})$. From the above conclusion (ii), we have $\lim_{i \rightarrow \infty} \|x_t^{s_i+1} - x_{t-1}^{s_i+1}\| = 0$ and $\lim_{i \rightarrow \infty} \|x_t^{s_i+1} - x_m^{s_i}\| = 0$. Then using the limit $s \rightarrow \infty$ in (41), we have

$$0 \in \nabla f(\bar{x}) + \partial g(\bar{x}).$$

Thus, we prove that \bar{x} is a stationary point of $\Phi(x)$. ■

Appendix E. The proof of Lemma 6

Proof From Lemma 5, we have $\{(H_t^s)_{t=0}^m\}_{s=1}^\infty$ is convergent, and

$$\lim_{s \rightarrow \infty} \|x_t^{s+1} - x_{t-1}^{s+1}\|^2 = 0, \text{ and } \lim_{s \rightarrow \infty} \|x_t^{s+1} - x_m^s\|^2 = 0.$$

These together with the definition of H_t^s implies that $\lim_{s \rightarrow \infty} \mathbb{E}[\Phi(x_t^s)]$. In the following, let $\lim_{s \rightarrow \infty} \mathbb{E}[\Phi(x_t^s)] = \zeta$.

Next, we will show that $\Phi \equiv \zeta$ on Ω . Since Ω denotes the set of accumulation points of the sequence $\{(x_t^s)_{t=1}^m\}_{s=1}^\infty$, for any $\bar{x} \in \Omega$, there exists a convergent subsequence $\{x_t^{s_i}\}$ so that

$$\lim_{i \rightarrow \infty} x_t^{s_i} = \bar{x}, \text{ for } \forall t \in \{0, 1, 2, \dots, m\}. \quad (42)$$

Using the lower semi-continuity of $\Phi(x)$, thus, we have

$$\Phi(\bar{x}) \leq \liminf_{i \rightarrow \infty} \mathbb{E}[\Phi(x_t^{s_i})] = \zeta. \quad (43)$$

By the definition of $x_t^{s_i}$ as the minimizer of problem (4), we have

$$g(x_{t+1}^{s_i+1}) + \langle v_t^{s_i+1}, x_{t+1}^{s_i+1} - \bar{x} \rangle + \frac{1}{2\mu} \|x_{t+1}^{s_i+1} - y_t^{s_i+1}\|^2 \leq g(\bar{x}) + \frac{1}{2\mu} \|\bar{x} - y_t^{s_i+1}\|^2. \quad (44)$$

Adding $f(x_{t+1}^{s_i+1})$ to both sides of (44), we have

$$\begin{aligned} & f(x_{t+1}^{s_i+1}) + g(x_{t+1}^{s_i+1}) + \langle v_t^{s_i+1}, x_{t+1}^{s_i+1} - \bar{x} \rangle + \frac{1}{2\mu} \|x_{t+1}^{s_i+1} - y_t^{s_i+1}\|^2 \\ & \leq f(x_{t+1}^{s_i+1}) + g(\bar{x}) + \frac{1}{2\mu} \|\bar{x} - y_t^{s_i+1}\|^2. \end{aligned} \quad (45)$$

Using $y_t^{s_i+1} = x_t^{s_i+1} + \beta_t(x_t^{s_i+1} - x_{t-1}^{s_i+1})$, we have

$$\begin{aligned} \|x_{t+1}^{s_i+1} - y_t^{s_i+1}\| &= \|x_{t+1}^{s_i+1} - x_t^{s_i+1} - \beta_t(x_t^{s_i+1} - x_{t-1}^{s_i+1})\| \\ &\leq \|x_{t+1}^{s_i+1} - x_t^{s_i+1}\| + \bar{\beta} \|x_t^{s_i+1} - x_{t-1}^{s_i+1}\|. \end{aligned} \quad (46)$$

Thus, we have

$$\begin{aligned} \|\bar{x} - y_t^{s_i+1}\| &= \|\bar{x} - x_t^{s_i+1} + x_t^{s_i+1} - y_t^{s_i+1}\| \\ &\leq \|\bar{x} - x_t^{s_i+1}\| + \|x_t^{s_i+1} - y_t^{s_i+1}\| \\ &\stackrel{(i)}{\leq} \|\bar{x} - x_t^{s_i+1}\| + \|x_{t+1}^{s_i+1} - x_t^{s_i+1}\| + \bar{\beta} \|x_t^{s_i+1} - x_{t-1}^{s_i+1}\|, \end{aligned} \quad (47)$$

where the inequality (i) holds by (46).

Combining (42), (46) with (47), we have

$$\lim_{i \rightarrow \infty} \|x_{t+1}^{s_i+1} - y_t^{s_i+1}\| = 0 \text{ and } \lim_{i \rightarrow \infty} \|\bar{x} - y_t^{s_i+1}\| = 0. \quad (48)$$

Using (48) with $\mathbb{E}[v_t^{s_i+1}] = \nabla f(y_t^{s_i+1})$, then we have

$$\zeta = \lim_{i \rightarrow \infty} \sup \mathbb{E}[\Phi(x_{t+1}^{s_i+1})] \leq \Phi(\bar{x}). \quad (49)$$

Finally, combining (43) with (49), we have $\Phi(\bar{x}) = \lim_{i \rightarrow \infty} \mathbb{E}[\Phi(x_t^{s_i})] = \zeta$ for any $\bar{x} \in \Omega$. ■

Appendix F. The proof of Theorem 7

Proof We first consider $\mathbb{E}[\|\text{Prox}_{\mu g}(x_t^s - \mu \nabla f(x_t^s)) - x_t^s\|^2]$ as follows:

$$\begin{aligned}
& \mathbb{E}[\|\text{Prox}_{\mu g}(x_t^s - \mu \nabla f(x_t^s)) - x_t^s\|^2] \\
& \leq 4\|\text{Prox}_{\mu g}(x_t^s - \mu \nabla f(x_t^s)) - \text{Prox}_{\mu g}(y_t^s - \mu \nabla f(y_t^s))\|^2 + 4\|y_t^s - x_t^s\|^2 \\
& \quad + 4\mathbb{E}[\|\text{Prox}_{\mu g}(y_t^s - \mu \nabla f(y_t^s)) - \text{Prox}_{\mu g}(y_t^s - \mu v_t^s)\|^2] + 4\|\text{Prox}_{\mu g}(x_t^s - \mu v_t^s) - y_t^s\|^2 \\
& \stackrel{(i)}{\leq} 4\|x_t^s - y_t^s - \mu(\nabla f(x_t^s) - \nabla f(y_t^s))\|^2 + 4\mu^2\mathbb{E}[\|\nabla f(y_t^s) - v_t^s\|^2] + 4\|y_t^s - x_t^s\|^2 \\
& \quad + 4\|x_{t+1}^s - y_t^s\|^2 \\
& \stackrel{(ii)}{\leq} (12 + 8\mu^2)\|y_t^s - x_t^s\|^2 + \frac{2\mu^2 L^2}{b}\|x_t^s - x_m^{s-1}\|^2 + \frac{2\mu^2 L^2 \bar{\beta}^2}{b}\|x_t^s - x_{t-1}^s\|^2 \\
& \quad + 4\|x_{t+1}^s - y_t^s\|^2 \\
& \stackrel{(iii)}{\leq} (12 + 8\mu^2)\bar{\beta}^2\|x_t^s - x_{t-1}^s\|^2 + \frac{2\mu^2 L^2}{b}\|x_t^s - x_m^{s-1}\|^2 + \frac{2\mu^2 L^2 \bar{\beta}^2}{b}\|x_t^s - x_{t-1}^s\|^2 \\
& \quad + 8\|x_{t+1}^s - x_t^s\|^2 + 8\bar{\beta}^2\|x_t^s - x_{t-1}^s\|^2 \\
& = (20 + 8\mu^2 + \frac{2\mu^2 L^2}{b})\bar{\beta}^2\|x_t^s - x_{t-1}^s\|^2 + 8\|x_{t+1}^s - x_t^s\|^2 + \frac{2\mu^2 L^2}{b}\|x_t^s - x_m^{s-1}\|^2, \tag{50}
\end{aligned}$$

where the inequality (i) holds by the non-expansiveness property of the proximal operator; the inequality (ii) holds by Lemma 3, and the inequality (iii) holds by the definition of y_t^s .

Recall the conclusion (ii) of Lemma 4, we have for any $t \in \{1, 2, \dots, m\}$,

$$\lim_{s \rightarrow \infty} \|x_t^s - x_{t-1}^s\|^2 = 0, \text{ and } \lim_{s \rightarrow \infty} \|x_t^s - x_m^{s-1}\|^2 = 0. \tag{51}$$

Combining (50) with (51), we have $\lim_{s \rightarrow \infty} \mathbb{E}[\|\text{Prox}_{\mu g}(x_t^s - \mu \nabla f(x_t^s)) - x_t^s\|^2] = 0$.

Since $\{(H_t^s)_{t=0}^m\}_{s=1}^\infty$ is non-increasing from Lemma 4, we have $H_t^s \leq H_0^1 = \xi$ for all $s > 0$ and $t \in \{0, 1, 2, \dots, m\}$. By Assumption 2, there exist $\gamma > 0$ and a positive integer S such that for all $s \geq S$, we have

$$\text{dist}(x_t^s, \mathcal{X}) \leq \gamma \mathbb{E}[\|\text{Prox}_{\mu g}(x_t^s - \mu \nabla f(x_t^s)) - x_t^s\|]. \tag{52}$$

So this proves (i).

Using (13), for any $z \in \mathcal{X}$, we have

$$\begin{aligned}
\Phi(x_{t+1}^{s+1}) & \leq \Phi(z) + \frac{L^2}{4\tau b}\|x_t^{s+1} - x_m^s\|^2 + \frac{\bar{\beta}^2 L^2}{4\tau b}\|x_t^{s+1} - x_{t-1}^{s+1}\|^2 \\
& \quad + \left(\frac{1}{2\mu} + \frac{l}{2}\right)\|z - y_t^{s+1}\|^2, \\
& = \Phi(z) + \frac{L^2}{4\tau b}\|x_t^{s+1} - x_m^s\|^2 + \frac{\bar{\beta}^2 L^2}{4\tau b}\|x_t^{s+1} - x_{t-1}^{s+1}\|^2 \\
& \quad + \left(\frac{1}{2\mu} + \frac{l}{2}\right)\|z - x_t^{s+1} + x_t^{s+1} - y_t^{s+1}\|^2, \\
& \leq \Phi(z) + \frac{L^2}{4\tau b}\|x_t^{s+1} - x_m^s\|^2 + \frac{\bar{\beta}^2 L^2}{4\tau b}\|x_t^{s+1} - x_{t-1}^{s+1}\|^2 \\
& \quad + \left(\frac{1}{\mu} + l\right)\|z - x_t^{s+1}\|^2 + \left(\frac{1}{\mu} + l\right)\|x_t^{s+1} - y_t^{s+1}\|^2. \tag{53}
\end{aligned}$$

Let $z = \hat{x}_t^{s+1} \in \mathcal{X}$, so that $\text{dist}(x_t^{s+1}, \mathcal{X}) = \|\hat{x}_t^{s+1} - x_t^{s+1}\|$, then we have

$$\begin{aligned} \Phi(x_{t+1}^{s+1}) - \Phi(\hat{x}_t^{s+1}) &\leq \frac{L^2}{4\tau b} \|x_t^{s+1} - x_m^s\|^2 + \frac{\bar{\beta}^2 L^2}{4\tau b} \|x_t^{s+1} - x_{t-1}^{s+1}\|^2 \\ &\quad + \left(\frac{1}{\mu} + l\right) \text{dist}(x_t^{s+1}, \mathcal{X})^2 + \left(\frac{1}{\mu} + l\right) \|x_t^{s+1} - y_t^{s+1}\|^2. \end{aligned} \quad (54)$$

By (53) and (54), we have $\lim_{s \rightarrow \infty} \|\hat{x}_{t+1}^{s+1} - x_{t+1}^{s+1}\|^2 = 0$. These together with the fact that

$$\begin{aligned} \|\hat{x}_{t+1}^{s+1} - \hat{x}_t^{s+1}\|^2 &= \|\hat{x}_{t+1}^{s+1} - x_{t+1}^{s+1} + x_{t+1}^{s+1} - x_t^{s+1} + x_t^{s+1} - \hat{x}_t^{s+1}\|^2 \\ &\leq 3\|\hat{x}_{t+1}^{s+1} - x_{t+1}^{s+1}\|^2 + 3\|x_{t+1}^{s+1} - x_t^{s+1}\|^2 + 3\|x_t^{s+1} - \hat{x}_t^{s+1}\|^2, \end{aligned}$$

implies that $\lim_{s \rightarrow \infty} \|\hat{x}_{t+1}^{s+1} - \hat{x}_t^{s+1}\|^2 = 0$. By Assumption 3, there exist a constant ζ so that $\Phi(\hat{x}_t^s) \equiv \zeta$ for all sufficiently large s and any $t \in \{0, 1, 2, \dots, m\}$. Thus, we have

$$\begin{aligned} \Phi(x_{t+1}^{s+1}) - \zeta &\leq \frac{L^2}{4\tau b} \|x_t^{s+1} - x_m^s\|^2 + \frac{\bar{\beta}^2 L^2}{4\tau b} \|x_t^{s+1} - x_{t-1}^{s+1}\|^2 \\ &\quad + \left(\frac{1}{\mu} + l\right) \text{dist}(x_t^{s+1}, \mathcal{X})^2 + \left(\frac{1}{\mu} + l\right) \|x_t^{s+1} - y_t^{s+1}\|^2, \end{aligned} \quad (55)$$

for all sufficiently large s and any $t \in \{0, 1, 2, \dots, m\}$.

Since \hat{x}_t^{s+1} is a stationary point of the problem (1), so that $-\nabla f(\hat{x}_t^{s+1}) \in \partial g(\hat{x}_t^{s+1})$. Due to the convexity of function $g(x)$, then we have

$$g(x_t^{s+1}) \geq g(\hat{x}_t^{s+1}) + \langle -\nabla f(\hat{x}_t^{s+1}), x_t^{s+1} - \hat{x}_t^{s+1} \rangle. \quad (56)$$

Using the definition of H_t^{s+1} , we have, for all sufficiently large s and any $t \in \{0, 1, 2, \dots, m\}$,

$$\begin{aligned} \zeta - H_t^{s+1} &= \Phi(\hat{x}_t^{s+1}) - \Phi(x_t^{s+1}) - \alpha(\|x_t^{s+1} - x_m^s\|^2 + \|x_t^{s+1} - x_{t-1}^{s+1}\|^2) \\ &= f(\hat{x}_t^{s+1}) + g(\hat{x}_t^{s+1}) - f(x_t^{s+1}) - g(x_t^{s+1}) - \alpha(\|x_t^{s+1} - x_m^s\|^2 + \|x_t^{s+1} - x_{t-1}^{s+1}\|^2) \\ &\leq f(\hat{x}_t^{s+1}) - f(x_t^{s+1}) + \langle -\nabla f(\hat{x}_t^{s+1}), \hat{x}_t^{s+1} - x_t^{s+1} \rangle - \alpha(\|x_t^{s+1} - x_m^s\|^2 + \|x_t^{s+1} - x_{t-1}^{s+1}\|^2) \\ &\stackrel{(i)}{\leq} \frac{L}{2} \|\hat{x}_t^{s+1} - x_t^{s+1}\|^2 - \alpha(\|x_t^{s+1} - x_m^s\|^2 + \|x_t^{s+1} - x_{t-1}^{s+1}\|^2), \end{aligned} \quad (57)$$

where the inequality (i) holds by the Lipschitz continuity of $-\nabla f(x)$.

From the conclusion (i), we have $\lim_{s \rightarrow \infty} \text{dist}(x_t^{s+1}, \mathcal{X}) = \lim_{s \rightarrow \infty} \|\hat{x}_t^{s+1} - x_t^{s+1}\| = 0$. This together with (51) implies the fact that

$$\zeta \leq \lim_{s \rightarrow \infty} H_t^{s+1} = \inf_s H_t^{s+1}, \quad (58)$$

where the equality holds by the sequence $\{(H_t^s)_{t=0}^m\}_{s=1}^\infty$ is non-increasing.

Next, combining (50), (52) with (55), for all sufficiently large s and any $t \in \{0, 1, 2, \dots, m\}$, we have

$$\begin{aligned} \Phi(x_{t+1}^{s+1}) - \zeta &\leq \frac{L^2}{4\tau b} \|x_t^{s+1} - x_m^s\|^2 + \frac{\bar{\beta}^2 L^2}{4\tau b} \|x_t^{s+1} - x_{t-1}^{s+1}\|^2 + \left(\frac{1}{\mu} + l\right) \text{dist}(x_t^s, \mathcal{X})^2 + \left(\frac{1}{\mu} + l\right) \|x_t^{s+1} - y_t^{s+1}\|^2 \\ &\stackrel{(i)}{\leq} \gamma \left(\frac{1}{\mu} + l\right) \left[(20 + 8\mu^2 + \frac{2\mu^2 L^2}{b}) \bar{\beta}^2 \|x_t^{s+1} - x_{t-1}^{s+1}\|^2 + 8\|x_{t+1}^{s+1} - x_t^{s+1}\|^2 + \frac{2\mu^2 L^2}{b} \|x_t^{s+1} - x_m^s\|^2\right] \\ &\quad + \bar{\beta}^2 \left(\frac{1}{\mu} + l\right) \|x_t^{s+1} - x_{t-1}^{s+1}\|^2 + \frac{L^2}{4\tau b} \|x_t^{s+1} - x_m^s\|^2 + \frac{\bar{\beta}^2 L^2}{4\tau b} \|x_t^{s+1} - x_{t-1}^{s+1}\|^2 \\ &\leq C(\|x_{t+1}^{s+1} - x_t^{s+1}\|^2 + \|x_t^{s+1} - x_{t-1}^{s+1}\|^2 + \|x_t^{s+1} - x_m^s\|^2) \end{aligned} \quad (59)$$

for some positive constant C . Using (59) with the definition of H_t^s , we have

$$0 \leq H_{t+1}^{s+1} - \zeta \leq \eta(\|x_{t+1}^{s+1} - x_t^{s+1}\|^2 + \|x_t^{s+1} - x_{t-1}^{s+1}\|^2 + \|x_t^{s+1} - x_m^s\|^2), \quad (60)$$

where $\eta = C + \alpha$, and the first inequality holds by (58). Then, by (14), we can obtain

$$(H_{t+1}^{s+1} - \zeta) - (H_t^{s+1} - \zeta) \leq -\kappa(\|x_{t+1}^{s+1} - x_t^{s+1}\|^2 + \|x_t^{s+1} - x_{t-1}^{s+1}\|^2 + \|x_t^{s+1} - x_m^s\|^2). \quad (61)$$

Finally, combining (60) with (61), we have

$$(H_{t+1}^{s+1} - \zeta) - (H_t^{s+1} - \zeta) \leq -\frac{\kappa}{\eta}(H_{t+1}^{s+1} - \zeta). \quad (62)$$

By (62), thus, we have

$$0 \leq H_{t+1}^{s+1} - \zeta \leq \frac{1}{1 + \frac{\kappa}{\eta}}(H_t^{s+1} - \zeta) \leq \frac{1}{(1 + \frac{\kappa}{\eta})^{m+1}}(H_{t+1}^s - \zeta).$$

This implies that the sequence $\{(H_t^s)_{t=0}^m\}_{s=1}^\infty$ is Q -linearly convergent. ■

References

- Allen-Zhu Z, Hazan E. Variance reduction for faster non-convex optimization. arXiv preprint arXiv:1603.05643, 2016.
- Aravkin A, Davis D. A SMART Stochastic Algorithm for Nonconvex Optimization with Applications to Robust Machine Learning. arXiv preprint arXiv:1610.01101, 2016.
- Ben-Israel A, Mond B. What is invexity?. The ANZIAM Journal, 1986, 28(1): 1–9.
- Léon Bottou. Stochastic learning. In Advanced lectures on machine learning, pages:146–168, Springer, 2004.
- Bousquet O, Bottou L. The tradeoffs of large scale learning. Advances in neural information processing systems. 2008: 161–168.
- Defazio A, Bach F, Lacoste-Julien S. SAGA: A fast incremental gradient method with support for non-strongly convex composite objectives. Advances in Neural Information Processing Systems, pages:1646–1654, 2014.
- Drusvyatskiy D, Lewis A S. Error bounds, quadratic growth, and linear convergence of proximal methods. arXiv preprint arXiv:1602.06661, 2016.
- Ghadimi S, Lan G. Accelerated gradient methods for nonconvex nonlinear and stochastic programming. Mathematical Programming, 156(1-2):59–99, 2016.
- Ghadimi S, Lan G, Zhang H. Mini-batch stochastic approximation methods for nonconvex stochastic composite optimization. Mathematical Programming, 155(1-2):267–305, 2016.

- Huang F, Chen S, Lu Z. Stochastic Alternating Direction Method of Multipliers with Variance Reduction for Nonconvex Optimization. arXiv preprint arXiv:1610.02758, 2016.
- Johnson R, Zhang T. Accelerating stochastic gradient descent using predictive variance reduction. In Advances in Neural Information Processing Systems, pages:315–323, 2013.
- Karimi H, Nutini J, Schmidt M. Linear convergence of gradient and proximal-gradient methods under the polyak-Łojasiewicz condition. Joint European Conference on Machine Learning and Knowledge Discovery in Databases. Springer International Publishing, 2016: 795–811.
- LeCun Y, Bengio Y, Hinton G. Deep learning. Nature, 521(7553):436–444, 2015.
- Luo Z Q, Tseng P. Error bounds and convergence analysis of feasible descent methods: a general approach. Annals of Operations Research, 1993, 46(1): 157–178.
- Yurii Nesterov. Introductory Lectures on Convex Programming Volume I: Basic course. Kluwer Boston, 2004.
- Nitanda A. Stochastic proximal gradient descent with acceleration techniques. Advances in Neural Information Processing Systems. 2014: 1574–1582.
- Polyak B T. Gradient methods for the minimisation of functionals. USSR Computational Mathematics and Mathematical Physics, 1963, 3(4): 864–878.
- Sashank J Reddi, Ahmed Hefny, Suvrit Sra, Barnabás Póczós and Alex Smola. Stochastic Variance Reduction for Nonconvex Optimization. arXiv preprint arXiv:1603.06160, 2016.
- Sashank J Reddi, Suvrit Sra, Barnabas Poczos and Alex Smola. Fast Stochastic Methods for Nonsmooth Nonconvex Optimization. arXiv preprint arXiv:1605.0690, 2016.
- Nicolas L Roux, Mark Schmidt and Francis R Bach. A stochastic gradient method with an exponential convergence rate for finite training sets. In Advances in Neural Information Processing Systems, pages:2663–2671, 2012.
- Shai Shalev-Shwartz and Tong Zhang. Stochastic dual coordinate ascent methods for regularized loss minimization. Journal of Machine Learning Research, 14:567–599, 2013.
- Tseng P. Approximation accuracy, gradient methods, and error bound for structured convex optimization. Mathematical Programming, 2010, 125(2): 263–295.
- Vapnik V. The nature of statistical learning theory. Springer Science & Business Media, 2013.
- Wen B, Chen X, Pong T K. Linear convergence of proximal gradient algorithm with extrapolation for a class of nonconvex nonsmooth minimization problems. SIAM Journal on Optimization, 2017, 27(1): 124–145.
- Xiao L, Zhang T. A proximal stochastic gradient method with progressive variance reduction. SIAM Journal on Optimization, 24(4):2057–2075, 2014.

Zhou Z, So A M C. A unified approach to error bounds for structured convex optimization problems. arXiv preprint arXiv:1512.03518, 2015.